

# CSEP561 – Jumble

---

David Wetherall

[djw@cs.washington.edu](mailto:djw@cs.washington.edu)

# mchat

---

- Test away on IP: 230.0.0.1
- Port:
  - First, find one (1) other team, pick a random port (between 10K and 20K, say), and test for interoperability.
  - If it works you survived round #1
  - Merge with a neighboring team for round #2
  - Repeat. Who will get to what round?
  - If/when there is a failure, how do we know who is at fault?
  
  - Afterwards, we'll all try port 4446!

# SST – question on DNS

---

- Benefits:
  - Can seamlessly support short and long replies (e.g., queries and zone transfers) with the same API
  - Burden of reliability is moved from resolver code to SST
- Costs
  - DNS servers can no longer be stateless. This is a big loss.
  - Latency for queries won't improve; in fact it requires SST to be carefully optimized so it does no harm (e.g., query bundled with connection setup) and impact is hard to gauge (e.g., graphs are for warm connections)

# SST – question on HTTP

---

- This is a motivating scenario for developing SST
- Benefits:
  - Much simpler HTTP code for top-of-the line performance (pipelining).
  - Performance may improve, but only if it wasn't very good already.
- Issues:
  - Compatibility. Need to change both ends. Oh well.

# SST – question on evaluation

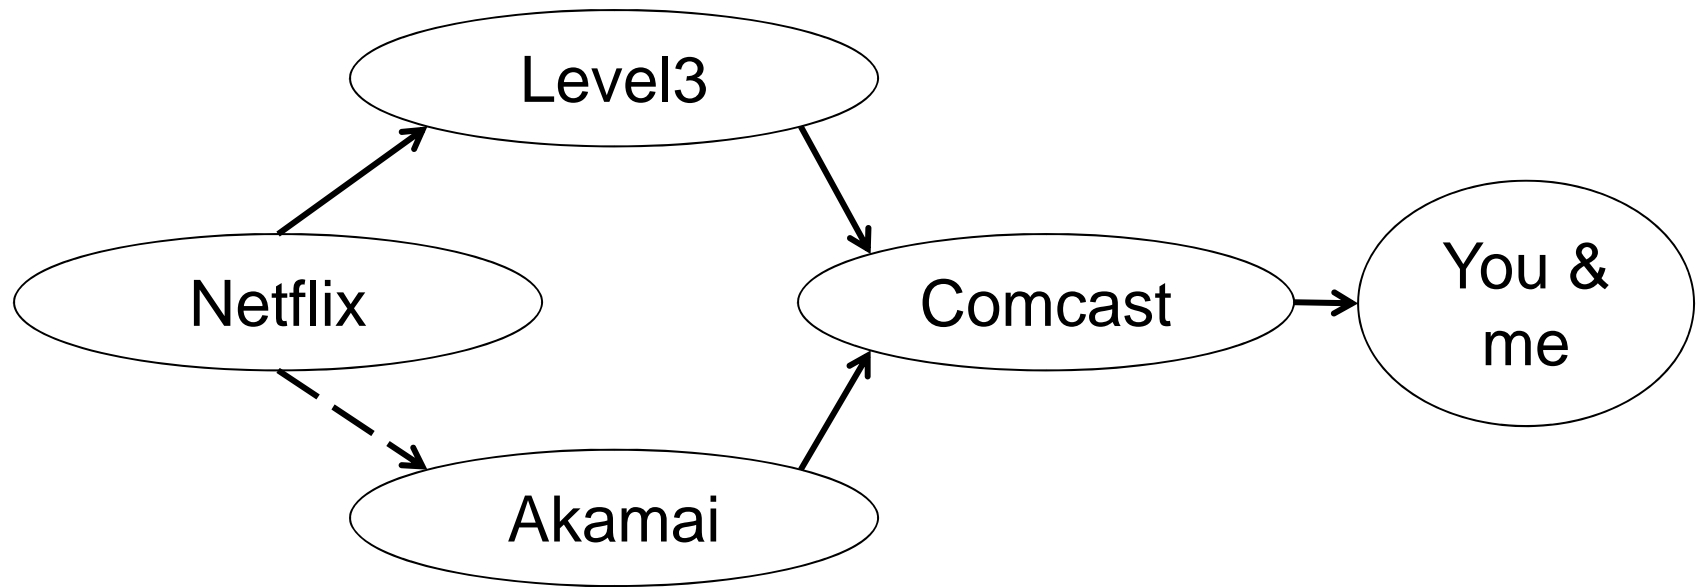
---

- It's performance-centric
  - On the whole does a good job, but ...
- Two key issues are likely to be:
  - Application writing experience: is it a good/better API?
  - How can we use it in practice (out of scope – later work)

# Comcast, Level3, NetFlix, and Akamai

---

- This is the flow of content. What's the flow of money?

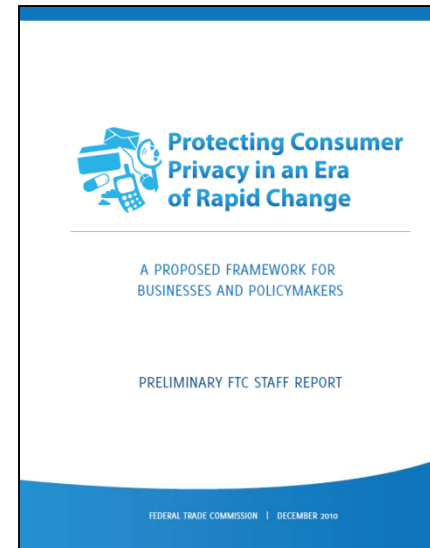


# Web Tracking

---

- Web analytics is perhaps a \$1B market (Google Analytics, Omniture)
- Try something like Ghostery to see how much is going on
  - [www.ghostery.com](http://www.ghostery.com)
- Some tracking is aggressive (individual profiles across independent websites)
- Yesterday the FTC recommended a “do not track” registry

Google Analytics



# How Tracking works

---

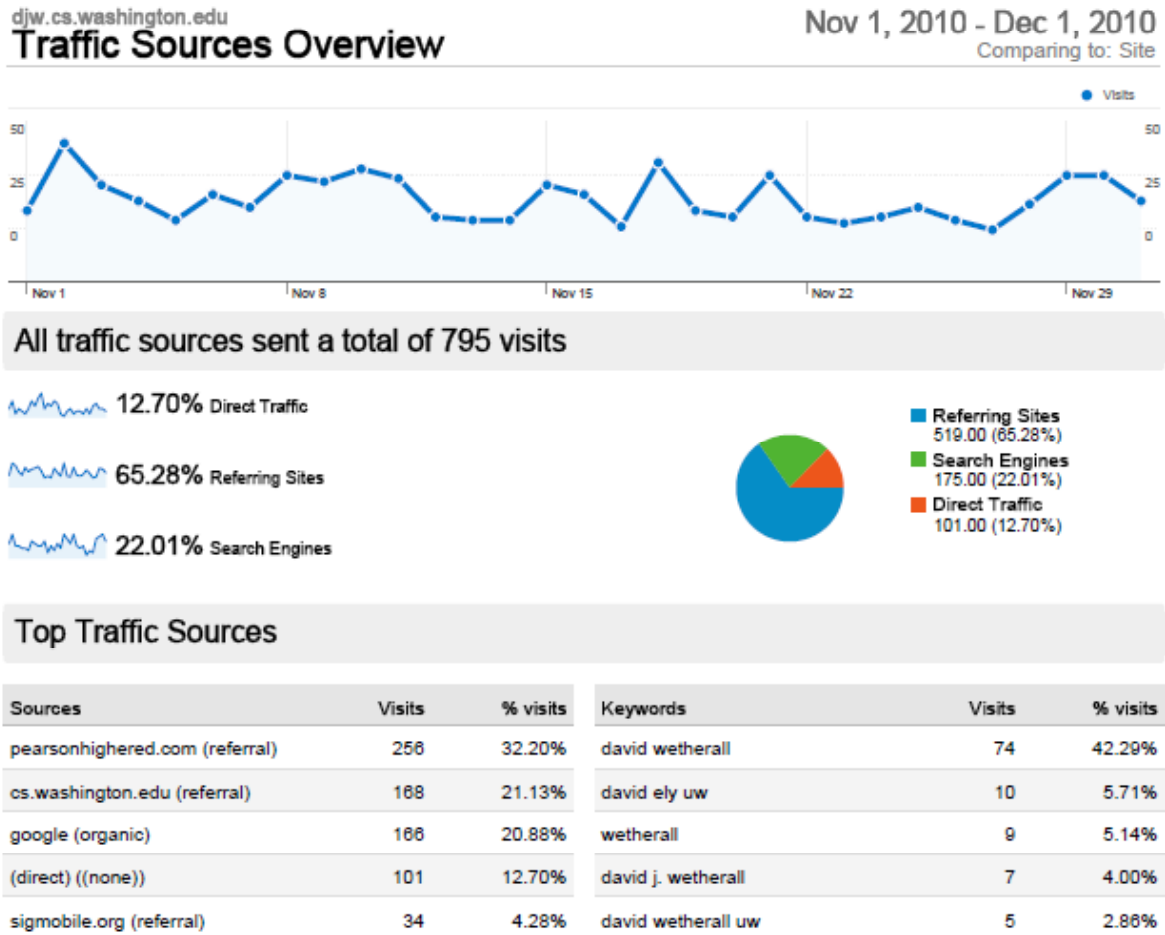
- Plant a small script in the pages you want to track:

```
<script type="text/javascript"> var gaJsHost = (("https:" ==  
document.location.protocol) ? "https://ssl." : "http://www.");  
document.write(unescape("%3Cscript src=" + gaJsHost + "google-  
analytics.com/ga.js" type='text/javascript'%3E%3C/script%3E")); </script> <script  
type="text/javascript"> try { var pageTracker = _gat._getTracker("UA-11438709-1");  
pageTracker._trackPageview(); } catch(err) {}</script>
```

- Loading page causes browser to make an invisible HTTP request of tracker; tracker learns what IP is loading which page, and what previous Web page they visited
- Tracker can send back a unique cookie. Subsequent requests will send the cookie back to the tracker, uniquely identifying the browser for profiling



# Google Analytics example



# True/False

---

- TCP/IP is the network layer protocol of the Internet
- The key mechanism to combine protocols is layering
- Lower layer data is encapsulated by higher layers
- When processing protocol data units, the next higher layer is found with a demux key
- The key disadvantage of layering is inefficiency
- The E2E argument says that functionality should be implemented at a higher layer when it improves performance

# True/False

---

- Bandwidth is measured in Hz and bps
- Fiber has very low attenuation but limited bandwidth
- The capacity of a channel is limited by its bandwidth and signal-to-noise ratio
- NRZ is a form of passband modulation
- Passband transmission requires a carrier to be modulated
- 4B/5B encoding is helpful for scrambling a signal

# True/False

---

- Message latency depends on both propagation delay and transmission delay
- The bandwidth-delay product tells us how quickly a remote receiver can react
- Wired links can be engineered for the best case; wireless links must be engineered for the worst case
- Important properties of a link are its rate, delay, error rate, and whether it is broadcast

# True/False

---

- Error detection codes typically need less overhead than error correction codes
- The Hamming distance of a code tells us how many bit errors can reliably be detected, but not corrected
- Checksums catch more errors than CRCs of the same size
- The BER is sufficient for designing an error correction scheme
- Interleaving is a useful technique for tolerating burst errors

# True/False

---

- Sequence numbers must be added to the sliding window protocol when the window is larger than one packet
- Stop-and-wait is efficient when the bandwidth-delay product is small
- Another name for a retransmission scheme is Forward Error Correction (FEC)
- TCP and IP use FEC for reliability
- Statistical multiplexing is much better suited to networks than static multiplexing because traffic is bursty

# True/False

---

- Randomized access protocols like classic Ethernet suffer from contention at high load
- Contention-free designs like rings are most efficient at low load
- Collision detection works best for wired links and when the bandwidth-delay product is small
- Binary Exponential Backoff adjusts the aggressiveness of medium access to match the number of contending nodes
- LANs are plug-and-play because of the “backward learning” algorithm
- Bridges work at the IP layer

# True/False

---

- The bridge spanning tree algorithm computes one spanning tree for each source
- Routing is a local process, forwarding is a global one
- The main types of routing are link-state and shortest path routing
- A common method of controlling which paths are used is by setting link weights and using “lowest cost” paths
- Link-state routing suffers from poor convergence compared to Distance-vector routing
- The sink tree for a network and given sink is unique
- Link-state is more scalable than distance vector



# True/False

---

- Routing today typically adapts to changing traffic levels automatically
- Besides unicast, other delivery models are broadcast, multicast, and anycast
- Core-based multicast trees are more complex than per-sender multicast trees, but they deliver better performance
- Core-based trees are more scalable than per-sender trees because they require nodes to keep less state
- Anycast is used in the Internet today

# True/False

---

- The key problems of internetworking are heterogeneity and scale
- The point of path MTU discovery is to maintain connectivity despite link heterogeneity
- IP addresses can be assigned automatically with ARP
- The error control protocol that is used with IP is DHCP
- A key design feature of Mobile IP is that routers are unchanged and unaware of mobile hosts
- A costs of having routers unaware of the location of mobile hosts is lower quality routes

# True/False

---

- BGP is closer to link-state than distance-vector routing
- The key feature of BGP that distinguishes it from intradomain routing is routing policy
- IP addresses are allocated in blocks called classes
- Mechanisms to make IP routing scale include prefixes, subnets, aggregation, and AS-level paths
- Routing with BGP is a two-level hierarchy: find the right AS, then find the right IP